Gibbsian On-Line Distributed Content Caching Strategy for Cellular Networks

Arpan Chattopadhyay, Bartłomiej Błaszczyszyn

arXiv:1610.02318v1 [cs.NI] 7 Oct 2016

***Abstract*—In this paper, we develop Gibbs sampling based techniques for learning the optimal placement of contents in a cellular network. We consider the situation where a finite collection of base stations are scattered on the plane, each covering a cell (possibly overlapping with other cells). Mobile users request for downloads from a finite set of contents according to some popularity distribution which may be known or unknown to the base stations. Each base station has a fixed memory space that can store only a strict subset of the contents at a time; hence, if a user requests for a content that is not stored at any of its serving base stations, the content has to be downloaded from the backhaul. Hence, we consider the problem of optimal content placement which minimizes the rate of download from the backhaul, or equivalently maximize the cache hit rate. It is known that, when multiple cells can overlap with one another (e.g., under dense deployment of base stations in small cell networks), it is not optimal to place the most popular contents in each base station. However, the optimal content placement problem is NP-complete. Using ideas of Gibbs sampling, we propose simple sequential content update rules that decide whether to store a content at a base station (if required from the base station) and which content has to be removed from the corresponding cache, based on the knowledge of contents stored in its neighbouring base stations. The update rule is shown to be asymptotically converging to the optimal content placement for all nodes under the knowledge of content popularity. Next, we extend the algorithm to address the situation where content popularities and cell topology are initially unknown, but are estimated as new requests arrive to the base stations; we show that our algorithm working with the running estimates of content popularities and cell topology also converges asymptotically to the optimal content placement. Finally, we demonstrate via a simple numerical computation the improvement in cache hit rate compared to most popular content placement and independent content placement strategies in the caches.**

***Index Terms*—Content caching, on-line cache update, cellular network, hit rate maximization, Gibbs sampling.**

I. Introduction

The proliferation of smartphones and tablets equipped with 3G and 4G connectivity and the fast growing demand for downloading multimedia files have resulted in severe overload in the internet backhaul, and it is expected to be worse with the advent of 5G in near future. Recent idea of densifying cellular networks will improve wireless throughput, but this will eventually push the backhaul bandwidth to its limit. In order to alleviate this problem, the idea of caching popular multimedia contents has recently been proposed. Given the fact that the popular contents are requested many times which results in network congestion, one way to reduce the congestion is to cache the popular contents at various intermediate nodes in the network. In case of cellular network, this requires adding physical memory to base stations (BSs): macro, micro, nano and pico. This has several advantages: (i) Caching contents at base stations reduce backhaul load. (ii) Caching reduces delay in fetching the content, thereby reducing the multimedia playback time. (iii) Caching will allow the end user to download a lower quality content in case his channel quality or bad or in case he wants to control his total amount of download.

Under dense placement of base stations, it is often the case that the *cells* (a cell is defined to be a region around a BS where the user is able to get sufficient downlink data rate from the BS) of different BSs might overlap with each other in an arbitrary manner (see [1]). Hence, if a user is covered by multiple BSs, he has the option to download a content from any one of the serving BSs. This gives rise to the problem of optimal content placement in the caches of cellular BSs (see [2], [3]); the trade-off is that ideally the caching strategy should avoid placing the same content (even if it has very high popularity) in two BSs whose cells have a significant overlap, while it is not desirable for the non-overlapped region. Optimal content placement under such situation requires global knowledge of base station locations and cell topologies, and solving the optimization problem requires intensive computation. In order to tackle these problems, we develop sequential cache update algorithms (motivated by Gibbs sampling) that asymptotically lead to optimal content placement, where each base station updates its contents only when a new content is downloaded from the backhaul to meet a user request, and this update is made solely based on the knowledge of the neighbouring BSs whose cells have nonzero intersection with the cell of the BS under consideration. The results are also extended to the case where the content popularities and cell topology are unknown initially and are learnt over time as new content requests arrive to the base stations. Simple numerical computation demonstrates the improvement in cache hit rate using Gibbs sampling technique for cache update, compared to most popular content placement and independent content placement strategies in the caches.

A. Related Work

There have been considerable amount of work in the literature dedicated to cellular caching. Benefits and challenges for caching in 5G networks have been described in [4]. The authors of [5] have developed a method to analyze the

Arpan Chattopadhyay and Bartłomiej Błaszczyszyn are with Inria/ENS, Paris, France. Email: arpan.chattopadhyay@inria.fr, Bartek.Blaszczyszyn@ens.fr

performance of caches (isolated or networked), and shown that placing the most popular subset of contents in each cache is not optimal in case of interconnected caches. The paper [3] deals with optimal content placement in wireless caches given BS-user association. The authors of [6] have addressed the problem of optimal content placement under user mobility. The authors of [2] have proposed a randomized content placement scheme in cellular BS caches in order to maximize cache hit rate, but their scheme assumes that the contents are placed independently across the caches, which is obviously suboptimal. This work was later extended to the case of heterogeneous networks in [7]. The authors of [8] have again considered independent probabilistic caching in a random heterogeneous network. The authors of [9] have addressed the problem of cache miss minimization in a random network setting. The authors of [10] have studied the problem of distributed caching in ultra-dense wireless small cell networks using mean field games; however, this formulation requires us to take base station density to infinity (which may not be true in practice), and it does not provide any guarantee on the optimality of this caching strategy. The paper [11] proposes a pricing based scheme for jointly assigning content requests to cellular BSs and updating the cellular caches; but this paper focuses on certain cost minimization instead of hit rate maximization, and it is optimal only when we can represent the data by very large number of chunks which can be used in employing rateless code.

The authors of [12] and [13] propose learning schemes for unknown time-varying popularity of contents, but their scheme does not have theoretical guarantee of convergence to the optimal content placement *across the network when cells of different BSs overlap with each other*. The paper [14] establishes that, when popularity is dynamic, any scheme that separates content popularity estimation and cache update (i.e., control) phases is strictly order-wise suboptimal in terms of hit rate.

Contrary to the prior literature, our current paper provides theoretical guarantee on convergence of the cache update scheme in a cellular network to the optimal placement (for the first time to the best of our knowledge), without the help of any centralized decision-maker. A weaker version of the results also holds when popularities are unknown initially and are learnt over time using the information of request arrivals in the base stations.

B. Organization and Our Contribution

The rest of the paper is organized as follows.

- The system model has been defined in Section II.
- In Section III, we propose an update scheme for the caches based on the knowledge of the contents cached in neighbouring BSs. The update scheme is based on Gibbs sampling techniques, and cache updates are made only when new content requests arrive. The scheme asymptotically converges to a near-optimal content placement in the network, since the scheme is proposed for a finite "inverse temperature" (to be defined later). We prove convergence of the proposed scheme. To the best of our knowledge, such a scheme has never been used in the context of caching in cellular network.
- In Section IV, we discuss how to slowly increase the inverse temperature to ∞ so that the near-optimal limiting solution in Section III actually converges to the globally optimal solution. We provide rigorous proof for the convergence of this scheme.
- In Section V, we discuss how to adapt the update schemes to the situation when unknown content popularities and cell topology are learnt over time as new content requests arrive to the BSs over time.
- In Section VI, we numerically demonstrate that the proposed Gibbs sampling approach has the potential to significantly improve the cache hit rate in cellular networks.
- Finally, we conclude in Section VII.

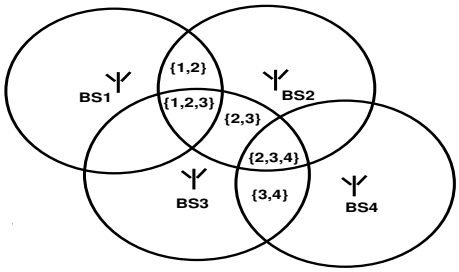

Fig. 1. A pictorial description of the base station (BS) coverage model. In the diagram, four BSs are shown with numbers $1, 2, 3, 4$. The circles correspond to the cells of the BSs. The region marked as $\{1, 2, 3\}$ has $s = \{1, 2, 3\}$; i.e., this region is called $\mathcal{C}(\{1, 2, 3\})$ and it is covered only by BSs $\{1, 2, 3\}$ and no other BS. Similar meaning applies to other regions.

II. System Model and Notation

A. Network Model

We consider a set $\mathcal{N} := \{1, 2, \cdots, N\}$ of base stations (BSs) on the two-dimensional Euclidean space. The set of points covered (downlink) by a BS constitute the *cell* of the corresponding BS.[1] We denote the cell of BS i ($1 \leq i \leq N$) by $\mathcal{C}_i$. Let us define $\mathcal{C} := \cup_{i=1}^{N} \mathcal{C}_i$. The area of any subset $\mathcal{A}$ of $\mathbb{R}^2$ is denoted by $|\mathcal{A}|$. We allow the cells of various BSs to have arbitrary and different *finite* areas. The cells of two BSs might have a nonzero intersection; any downlink mobile user located at such an intersection is covered by more than one BS. Let us denote by $2^{\mathcal{N}}$ the collection of all subsets of $\mathcal{N}$, and let s denote one such generic subset. Let us denote by $\mathcal{C}(s) := (\cap_{i \in s} \mathcal{C}_i) \cap (\cup_{i \notin s} \mathcal{C}_i)^c$ the region in $\mathcal{C}$ which is covered only by the BSs from the subset s.

See figure 1 for a better understanding of the cell model.

B. Content Request Process

Contents from a set $\mathcal{M} := \{1, 2, \cdots, M\}$ are requested by users located inside $\mathcal{C}$. We assume that each of these contents have the same size. Content i ($1 \leq i \leq M$) is

[1]This coverage could be signal to noise ration (SNR) based coverage where a point is covered by a BS if and only if the SNR at that point from the BS exceeds some threshold.

requested by users according to a homogeneous Poisson point process in space (inside $\mathcal{C}$) and time with intensity λ_i; this is the expected number of requests for content i per second per square meter inside $\mathcal{C}$. Let $\lambda := \sum_{i=1}^{M} \lambda_i$. Note that, $\frac{\lambda_i}{\lambda}$ denotes the probability that a content request is for content i; in other words, $\frac{\lambda_i}{\lambda}$ is the popularity of content i.

C. Content Caching at BSs

We assume that each BS can store K number of contents (where $K < M$). Let B denote a generic configuration of content placement in caches of the network. B is defined as a $M \times N$ matrix with $B_{i,j} = 1$ if content i is stored at the cache of BS j, and $B_{i,j} = 0$ otherwise. Note that, any feasible B must satisfy $\sum_{i=1}^{M} B_{i,j} = K$ for all $j \in \{1, 2, \cdots, N\}$; we rule out the possibility of $\sum_{i=1}^{M} B_{i,j} < K$ since that will be a waste of cache memory resources in BSs. Let us denote the set of all feasible configurations by $\mathcal{B}$. Clearly, the cardinality of $\mathcal{B}$ is $\binom{M}{K}^N$.

Apart from B, we will also use the symbol A for a generic configuration belonging to set $\mathcal{B}$.

D. Cache Hit Rate Maximization Problem

We assume that, whenever a new request for a content arrives, it is served by one BS covering that point and having the content in its cache; if a content request is served from the cache, we call the event as a *cache hit*. In case no covering BS has the content (i.e., no cache hit), the content needs to be downloaded by one of the covering BSs and served to the user (this will be explained later). The requests do not tolerate any delay; i.e., we do not consider the possibility of holding the requests in a queue and serving the content to users in batch once the content becomes available in a BS. Also, we assume infinite bandwidth available for all downlink transmissions; i.e., each content is assumed to be served instantaneously.[2]

Let the random variable H_B denote the number of cache hits in the entire network in unit time, under configuration B. We define the cache hit rate $h(B) = \mathbf{E}(H_B)$ where the expectation is over the randomness in the content request arrival process. Clearly,

$$h(B) = \sum_{s \in 2^{\mathcal{N}}} |\mathcal{C}(s)| \sum_{i=1}^{M} \lambda_i \mathbf{1}\{\sum_{j \in s} B_{i,j} \geq 1\} \tag{1}$$

In this paper, we are interested in finding an optimal configuration which achieves:

$$\sup_{B \in \mathcal{B}} h(B) \tag{2}$$

We know that (2) is is an optimization problem with $0-1$ integer variables, nonlinear objective function and linear constraints. This class of problems has been shown to be NP-complete (see [15]), and hence, we cannot expect any polynomial time algorithm to solve (2). Hence, in this section, we provide iterative, distributed cache update scheme that asymptotically solves the problem. However, since the algorithm is iterative, we cannot use the optimal configuration over infinite time horizon. Hence, we seek to design a randomized iterative cache update scheme which yields

$$\liminf_{T \to \infty} \frac{\int_0^T \mathbf{E}(h(R(\tau)))d\tau}{T} = \sup_{B \in \mathcal{B}} h(B) \tag{3}$$

where $R(\tau) \in \mathcal{B}$ is the configuration of all caches in the network at time τ. Our iterative scheme is randomized, which renders $R(\tau)$ a random variable; hence, we work with its expectation $\mathbf{E}$.

[2]This is a valid assumption when the downlink traffic in the network is light.

III. Cache Update via Basic Gibbs Sampling

Let us rewrite (1) as $h(B) = \sum_{j=1}^{N} h_j(B)$ where

$$h_j(B) = \sum_{i=1}^{M} \lambda_i \sum_{s \in 2^{\mathcal{N}}} \frac{|\mathcal{C}(s)| B_{i,j} \mathbf{1}\{j \in s\}}{\max\{1, \sum_{k \in s} B_{i,k}\}} \tag{4}$$

We call $h_j(B)$ to be the cache hit rate seen by BS j under configuration B. This will be the true cache hit rate seen by BS j under configuration B if a new content request is served by one covering BS chosen uniformly from the set of covering BSs having that content. Note that, if more than one covering BSs have that content, choice of the serving BS will not affect the hit rate; hence, we can safely assume uniform choosing of the serving BS.

In order to solve $\sup_{B \in \mathcal{B}} \sum_{j=1}^{N} h_j(B)$, we propose to employ Gibbs sampling techniques (see [16, Chapter 7]). Let us assume that each BS maintains a *virtual* cache capable of storing K contents. The broad idea is that one can update the virtual cache contents in an iterative fashion using Gibbs sampling. Whenever a content is requested from a BS not having the content in its physical (real) cache, the BS will download it from the backhaul and, at the same time, will decide to store it in the real cache depending on whether it is stored in its virtual cache or not.

We will update the virtual caches according to a stochastic iterative algorithm so that the steady state probability of configuration B becomes:

$$\pi_\beta(B) := \frac{e^{\beta h(B)}}{\sum_{B' \in \mathcal{B}} e^{\beta h(B')}} := \frac{e^{\beta h(B)}}{Z_\beta} \tag{5}$$

Here β is called the "inverse temperature" (motivated by literature from statistical Physics), and Z_β is called the *partition function*.

Note that, $\lim_{\beta \to \infty} \sum_{B \in \arg\max_{B' \in \mathcal{B}} h(B')} \frac{e^{\beta h(B)}}{\sum_{B' \in \mathcal{B}} e^{\beta h(B')}} = 1$. Hence, if we choose configuration B for all virtual caches with probability $\pi_\beta(B)$, then, for sufficiently large β, the chosen configuration will belong to $\arg\max_{B \in \mathcal{B}} h(B)$ with probability close to 1. If real cache configuration closely follows virtual cache configuration, we can achieve optimal cache hit rate for real caching system.

A. Gibbs sampling approach for "virtual" cache update

Let us consider discrete time instants $t = 0, 1, 2, \cdots$ when virtual cache contents are updated; this is different from the

continuous time τ used before. Let us denote the configuration in all virtual caches in the network after the t-th decision instant by $V(t)$, where $V(t) \in \mathcal{B}$.

The Gibbs sampling algorithm simulates a discrete-time Markov chain $V(t)$ on state space $\mathcal{B}$, whose stationary probability distribution is given by $\pi_\beta(B) = \frac{e^{\beta h(B)}}{Z_\beta}$.

Let us define the set of *neighbours* of BS j (**including BS** j) as $\Psi(j) := \{n : n \in \mathcal{N}, \mathcal{C}_j \cap \mathcal{C}_n \neq \emptyset\}$. Let us denote by $B_{\cdot,-j}$ the restriction of configuration B to all BSs except BS j, i.e., $B_{\cdot,-j}$ is obtained by deleting the j-th column of B. Let $\pi_\beta(\cdot|B_{\cdot,-j})$ denote the conditional distribution of the network-wide configuration conditioned on $B_{\cdot,-j}$, under the joint distribution $\pi_\beta(\cdot)$. Clearly, $\pi_\beta(A|B_{\cdot,-j}) = 0$ if $A_{\cdot,-j} \neq B_{\cdot,-j}$.

If $A_{\cdot,-j} = B_{\cdot,-j}$, then

$$\pi_\beta(A|B_{\cdot,-j}) = \frac{e^{\beta h(A)}}{\sum_{v_j \in \{0,1\}^M, ||v_j||_1 = K} e^{\beta h(v_j, B_{\cdot,-j})}} \quad (6)$$

where $||v_j||_1$ is the sum of all components of the vector v_j.

Note that, there is common factor $e^{\beta \sum_{n \notin \Psi(j)} h_n(B)}$ in both numerator and denominator of the expression in (6), since this term does not depend on the contents in the virtual cache in BS j. Hence, (6) can be further simplified as:

$$\pi_\beta(A|B_{\cdot,-j}) = \frac{e^{\beta \sum_{n \in \Psi(j)} h_n(A)}}{\sum_{v_j \in \{0,1\}^M, ||v_j||_1 = K} e^{\beta \sum_{n \in \Psi(j)} h_n(v_j, B_{\cdot,-j})}} \quad (7)$$

Now, let us define $h_n(A, s)$ to be the hit rate seen by BS n under configuration A due to the content requests generated from the region $\mathcal{C}(s)$. Clearly, $h_n(A) = \sum_{s \in 2^{\mathcal{N}}} h_n(A, s)$, since the hit rate at BS n under configuration A is equal to the sum of hit rates by requests generated from all possible segments $\{\mathcal{C}(s)\}_{s \in 2^{\mathcal{N}}}$. Now, note that, the term $e^{\beta \sum_{n \in \Psi(j)} \sum_{s: j \notin s} h_n(A,s)}$ is a common factor in the numerator and denominator of the expression in (7), since this factor does not depend on the contents in the virtual cache of BS j. Hence, when $A_{\cdot,-j} = B_{\cdot,-j}$, we can simplify (7) further as follows:

$$\pi_\beta(A|B_{\cdot,-j}) = \frac{e^{\beta \sum_{n \in \Psi(j), s \ni j} h_n(A,s)}}{\sum_{v_j \in \{0,1\}^M, ||v_j||_1 = K} e^{\beta \sum_{n \in \Psi(j), s \ni j} h_n(v_j, B_{\cdot,-j}, s)}} \quad (8)$$

where

$$h_n(A, s) = \sum_{i=1}^{M} \lambda_i \frac{|\mathcal{C}(s)| A_{i,n} \mathbf{1}\{n \in s\}}{\max\{1, \sum_{k \in s} A_{i,k}\}} \quad (9)$$

We now describe an algorithm for sequentially updating the network-wide virtual cache configuration $V(t)$.

Algorithm 1: Start with an arbitrary $V(0) \in \mathcal{B}$. At discrete time t, pick a node $j_t \in \mathcal{N}$ randomly having uniform distribution from $\mathcal{N}$. Then, update the contents in the virtual cache of BS j_t by picking up a network-wide virtual cache configuration $A \in \mathcal{B}$ with probability $\pi_\beta(A|V_{\cdot,-j_t}(t-1))$. Note that, only contents in the virtual cache of BS j_t are modified by this operation. □

Theorem 1: Under Algorithm 1, $\{V(t)\}_{t \geq 0}$ is a reversible Markov chain, and it achieves the steady-state probability distribution $\pi_\beta(B) = \frac{e^{\beta h(B)}}{Z_\beta}$.

Proof: The proof is standard, and it follows from the theory in [16, Chapter 7]). ■

Remark 1: In Algorithm 1, in order to make an update at time t, node j_t needs to know the contents of the virtual caches only from $\Psi(j_t)$. This requires information exchange in each slot. Such information exchange can use the backhaul network, and this does not exert much load on the backhaul since the actual contents are not exchanged via the backhaul in this process.

Remark 2: The denominator in the simplified sampling probability expression in (8) requires a summation over all possible virtual cache configurations in $\Psi(j_t)$. This allows the system to avoid the huge combinatorial problem of calculating Z_β which requires $\binom{M}{K}^N$ addition operations. The advantage will be even more visible if we consider the possibility of varying β with time or learning $\{\lambda_i\}_{1 \leq i \leq M}$ over time if they are not known; the optimization problem $\sup_{B \in \mathcal{B}} h(B)$ will change over time in this case, and it will require calculation of the partition function in each slot.

B. The real cache update scheme for fixed β

Now we propose a cache update scheme for the real caches present in the BSs. *Our scheme decides to store a content in the cache of a BS only when the content is requested from that BS*. This eliminates the necessity of any unnecessary download from the backhaul.

Let us consider content request arrivals at continuous time (denoted by τ again) to the BS. Let us recall that the virtual caches are updated only at discrete times $t = 0, 1, 2, \cdots$. We assume that these discrete time instants $t = 0, 1, 2, \cdots$ units are superimposed on the continuous time axis $\tau \geq 0$. Hence, $V(\tau)$ is defined to be equal to $V(t)$ for $\tau \in [t, t+1)$, where $t \in \mathbb{Z}_+$.

Let us consider an increasing sequence of positive real numbers (viewed as time durations) $T_1, T_2, T_3, \cdots$ such as $T_k \uparrow \infty$ as $k \to \infty$. Let $S_l := T_1 + T_2 + \cdots + T_l$. Let $\kappa(\tau) := \sup\{l \in \mathbb{Z}_+ : S_l \leq \tau\}$ and $\zeta(\tau) := S_{\kappa(\tau)}$.

The real cache update scheme is given as follows:

Algorithm 2: Start with some arbitrary $R(0) \in \mathcal{B}$. At time τ, if the request for content i arrives at BS j (either because no other covering BS has this content or because BS j has been chosen from among the covering BSs having content i), then BS j does the following:

If BS j has content i, it will serve that.

If BS j does not have content i, it serves the content by downloading from the backhaul. Then content i is stored in the real cache of BS j if and only if $V_{i,j}(\zeta(\tau)-) = 1$ (i.e., if content i was stored in the virtual cache of BS j at time $\zeta(\tau)-$). If the BS j decides to store content i then, in order to make room for the newly stored content i, any content $k \neq i$ such that $V_{k,j}(\zeta(\tau)-) = 0$ and $R_{k,j}(\tau) = 1$, is removed from the real cache of BS j. □

Remark 3: The idea behind taking $T_k \to \infty$ as $k \to \infty$ in Algorithm 2 is as follows. We know that $V(t)$ reaches the distribution $\pi_\beta(\cdot)$ as $t \to \infty$. As $k \to \infty$, the fraction of time spent during $\tau \in [S_k, S_{k+1})$ in copying the contents present in $V(S_k-)$ to real caches becomes negligible, and the real caches are allowed to operate larger and larger fraction of time under content distribution close to $\pi_\beta(\cdot)$.

Now we make the following assumption:

Assumption 1: $|\mathcal{C}_i \cap (\cup_{j\neq i}\mathcal{C}_j)^c| > 0$ for all $i \in \{1, 2, \cdots, N\}$.

Theorem 2: Under Assumption 1, Algorithm 1 and Algorithm 2, we have (for the real caches in all BSs):

$$\lim_{T\to\infty} \frac{\int_0^T \mathbf{P}(R(\tau) = B)d\tau}{T} = \frac{e^{\beta h(B)}}{\sum_{B'\in\mathcal{B}} e^{\beta h(B')}}$$

Proof: Fix a small $\epsilon > 0$. Under configuration B of the *virtual caches*, let us denote the total time T_B (a generic random variable) taken by the arrival process so that, for all possible pairs $(i, j) \in \{1, 2, \cdots, M\} \times \{1, 2, \cdots, N\}$ there is at least one arrival of content i to BS j if virtual configuration B suggests placing content i at BS j; clearly $\mathbf{E}(T_B) < \infty$, since we have made Assumption 1. Let us consider $l \in \mathbb{Z}_+$ large enough such that: (i) $\sum_{B\in\mathcal{B}} |\mathbf{P}(V(t) = B) - \pi_\beta(B)| < \epsilon$ for all integer $t \geq S_{l-1}$, (ii) $\mathbf{P}(T_B > \epsilon T_{l+1}) < \epsilon$.

Now,

$$\begin{aligned}
&\frac{\int_{S_l}^{S_{l+1}} \mathbf{P}(R(\tau) = B)d\tau}{T_{l+1}} \\
\geq\ & \frac{\mathbf{E}\int_{\min\{S_l+T_B', S_{l+1}\}}^{S_{l+1}} \mathbf{P}(R(\tau) = B)d\tau}{T_{l+1}} \\
\geq\ & \frac{(1-\epsilon)\int_{S_l+\epsilon T_{l+1}}^{S_{l+1}} \mathbf{P}(R(\tau) = B|T_B' \leq \epsilon T_{l+1})d\tau}{T_{l+1}} \\
=\ & \frac{(1-\epsilon)(T_{l+1} - \epsilon T_{l+1})\mathbf{P}(V(S_l-) = B)}{T_{l+1}} \\
\geq\ & (1-\epsilon)^2(\pi_\beta(B) - \epsilon)
\end{aligned} \tag{10}$$

where T_B' has the same distribution as T_B.

Hence,

$$\begin{aligned}
&\liminf_{T\to\infty} \frac{\int_0^T \mathbf{P}(R(\tau) = B)d\tau}{T} \\
=\ & \liminf_{T\to\infty} \frac{\int_{S_l}^T \mathbf{P}(R(\tau) = B)d\tau}{T - S_l} \\
\geq\ & (1-\epsilon)^2(\pi_\beta(B) - \epsilon)
\end{aligned}$$

Since $\epsilon > 0$ is arbitrarily small, we have:

$$\liminf_{T\to\infty} \frac{\int_0^T \mathbf{P}(R(\tau) = B)d\tau}{T} \geq \pi_\beta(B)$$

.

But, by Fatou's lemma,

$$\begin{aligned}
&\sum_{B\in\mathcal{B}} \liminf_{T\to\infty} \frac{\int_0^T \mathbf{P}(R(\tau) = B)d\tau}{T} \\
\leq\ & \liminf_{T\to\infty} \sum_{B\in\mathcal{B}} \frac{\int_0^T \mathbf{P}(R(\tau) = B)d\tau}{T} \\
=\ & 1
\end{aligned}$$

and $\sum_{B\in\mathcal{B}} \pi_\beta(B) = 1$. Hence, we must have $\liminf_{T\to\infty} \frac{\int_0^T \mathbf{P}(R(\tau)=B)d\tau}{T} = \pi_\beta(B)$ for all $B \in \mathcal{B}$.

On the other hand,

$$\limsup_{T\to\infty} \frac{\int_0^T \mathbf{P}(R(\tau) = B)d\tau}{T} = \limsup_{l\to\infty} \frac{\int_{S_l}^{S_{l+1}} \mathbf{P}(R(\tau) = B)d\tau}{T_{l+1}}$$

Now,

$$\begin{aligned}
&\frac{\int_{S_l}^{S_{l+1}} \mathbf{P}(R(\tau) = B)d\tau}{T_{l+1}} \\
\leq\ & \frac{\int_{S_l+\epsilon T_{l+1}}^{S_{l+1}} \mathbf{P}(R(\tau) = B)d\tau + \epsilon T_{l+1}}{T_{l+1}} \\
\leq\ & \frac{\int_{S_l+\epsilon T_{l+1}}^{S_{l+1}} \mathbf{P}(R(\tau) = B|T_B' \leq \epsilon T_{l+1})d\tau + 2\epsilon T_{l+1}}{T_{l+1}} \\
\leq\ & \frac{\int_{S_l+\epsilon T_{l+1}}^{S_{l+1}} \mathbf{P}(V(S_l-) = B)d\tau + 2\epsilon T_{l+1}}{T_{l+1}} \\
\leq\ & \pi_\beta(B) + 3\epsilon
\end{aligned}$$

where the second inequality follows from the fact that for $\tau \in [S_l + \epsilon T_{l+1}, S_{l+1})$, the following holds:

$$\begin{aligned}
&\mathbf{P}(R(\tau) = B) \\
\leq\ & \mathbf{P}(R(\tau) = B|T_B' \leq \epsilon T_{l+1}) \\
& + \mathbf{P}(T_B' > \epsilon T_{l+1})\mathbf{P}(R(\tau) = B|T_B' > \epsilon T_{l+1}) \\
\leq\ & \mathbf{P}(R(\tau) = B|T_B' \leq \epsilon T_{l+1}) + \epsilon
\end{aligned}$$

Since $\epsilon > 0$ is arbitrarily small, we can say that

$$\limsup_{T\to\infty} \frac{\int_0^T \mathbf{P}(R(\tau) = B)d\tau}{T} \leq \pi_\beta(B)$$

Hence, $\lim_{T\to\infty} \frac{\int_0^T \mathbf{P}(R(\tau)=B)d\tau}{T} = \pi_\beta(B)$. ■

Remark 4: Note that, Assumption 1 is very crucial in the proof of Theorem 2, because this assumption ensures that every BS gets content requests at some nonzero arrival rate, and hence can update its real cache at strictly positive rate.[3]

IV. Varying β to Reach Optimality

In this section, we discuss how to vary the inverse temperature β to infinity with time so that the Gibbs sampling

[3]If Assumption 1 is not satisfied, then one can still achieve near optimal hit rate in real caches. It is achieved under a scheme where a new content request is sent to any of its covering BSs with very small probability $\eta > 0$, and otherwise the request is sent to a covering BS having that content. Similar analysis as in this paper can show that the time-average expected hit rate under this scheme differs from the optimal hit rate $\max_{B\in\mathcal{B}} h(B)$ only by a small margin which goes to 0 as $\eta \downarrow 0$.

algorithm (used to update virtual caches) converges to the optimizer of (2). The intuition is that this, combined with Algorithm 2 used for real cache update, will achieve optimal time-average expected cache hit rate for problem (3).

Let us define

$$\Delta := \max_{B\in\mathcal{B}} h(B) - \min_{B^{'}\in\mathcal{B}} h(B^{'}) > 0$$

Algorithm 3: This algorithm is analogous to Algorithm 1 except that, at discrete time instant $t = lN, lN+1, \cdots, lN+N-1$, we use $\beta_t := \beta_0 \log(1+l)$ instead of fixed β, where $0 < \beta_0 < \infty$ is the initial inverse temperature satisfying $\beta_0 N\Delta < 1$ and $\beta_0 \max_{B\in\mathcal{B}} h(B) < 1$. □

Theorem 3: Under Algorithm 3 for virtual cache update, the discrete time non-homogeneous Markov chain $\{V(t)\}_{t\geq 0}$ is strongly ergodic, and the limiting distribution $\pi_{v,\infty}$ satisfies $\pi_{v,\infty}(\arg\max_{B\in\mathcal{B}} h(B)) = 1$.

Proof: In this proof, we will use the notion of weak and strong ergodicity of time-inhomogeneous Markov chains from [16, Chapter 6, Section 8]), which is provided in the appendix.

Fix $k = 0$. We will first show that the Markov chain $\{V(t)\}_{t\geq 0}$ in weakly ergodic.

Let us consider the transition probability matrix (t.p.m.) Q_l for the inhomogeneous Markov chain $\{Y(l)\}_{l\geq 0}$ (where $Y(l) := \{V(lN), V(lN+1), \cdots, V(lN+N-1)\}$) indexed by the period index l (when $\beta_l = \beta_0 \log(1+l)$). Then, the Dobrushin's ergodic coefficient $\delta(Q_l)$ is given by (see [16, Chapter 6, Section 7] for definition) $\delta(Q_l) = 1 - \inf_{\mathbf{B}^{'},\mathbf{B}^{''}\in\mathcal{B}^N} \sum_{\mathbf{B}\in\mathcal{B}^N} \min\{Q_l(\mathbf{B}^{'},\mathbf{B}), Q_l(\mathbf{B}^{''},\mathbf{B})\}$. The Markov chain is weakly ergodic if $\sum_{l=1}^{\infty}(1-\delta(Q_l)) = \infty$ (by [16, Chapter 6, Theorem 8.2]).

Now, with positive probability, virtual caches in all nodes are updated over a period of N slots. Hence, any $\mathbf{B} \in \mathcal{B}^N$ can be reached over a period of N slots, starting from any other $\mathbf{B}^{'} \in \mathcal{B}^N$. Note that, once a node j_t is chosen in Algorithm 1, the sampling probability for any set of contents in its virtual cache in a slot is lower bounded by $\frac{e^{-\beta_l\Delta}}{\binom{M}{K}}$. Hence, for independent sampling over N slots, we will always have $Q_l(\mathbf{B}^{'},\mathbf{B}) > \left(\frac{e^{-\beta_l\Delta}}{N\binom{M}{K}}\right)^N > 0$ for all pairs $\mathbf{B}^{'},\mathbf{B}$. Hence,

$$\begin{aligned}
&\sum_{l=0}^{\infty}(1-\delta(Q_l)) \\
=\ &\sum_{l=0}^{\infty} \inf_{\mathbf{B}^{'},\mathbf{B}^{''}\in\mathcal{B}^N} \sum_{\mathbf{B}\in\mathcal{B}^N} \min\{Q_l(\mathbf{B}^{'},\mathbf{B}), Q_l(\mathbf{B}^{''},\mathbf{B})\} \\
\geq\ &\sum_{l=0}^{\infty} \left(\frac{e^{-\beta_0\log(1+l)\times\Delta}}{N\binom{M}{K}}\right)^N \\
=\ &\frac{1}{(N\binom{M}{K})^N} \sum_{l=0}^{\infty} e^{-N\Delta\beta_0\log(1+l)} \\
=\ &\frac{1}{(N\binom{M}{K})^N} \sum_{l=0}^{\infty} \frac{1}{(1+l)^{\beta_0 N\Delta}} \\
=\ &\infty \qquad (11)
\end{aligned}$$

Here the last step follows from the fact that $\sum_{l=1}^{\infty}\frac{1}{l^a}$ diverges for $0 < a < 1$.

Hence, the Markov chain $\{Y(l)\}_{l\geq 0}$ is weakly ergodic. Hence, $\{V(t)\}_{t\geq 0}$ is also weakly ergodic Markov chain.

Now we will use [16, Chapter 6, Theorem 8.3] to prove strong ergodicity of $\{V(t)\}_{t\geq 0}$.

Let us denote the t.p.m. of $\{V(t)\}_{t\geq 0}$ at a specific time $t = T$ by $Q^{(T)}$ (a specific matrix). If the Markov chain $\{V(t)\}_{t\geq 0}$ is allowed to evolve up to infinite time with *fixed* t.p.m. $Q^{(T)}$, then we will get stationary distribution $\pi_{\beta_T}(B) = \frac{e^{\beta_T h(B)}}{Z_{\beta_T}}$. This satisfies Condition 8.9 of [16, Chapter 6, Theorem 8.3].

Now we will check Condition 8.10 of [16, Chapter 6, Theorem 8.3]. For any $B \in \arg\max_{B^{'}\in\mathcal{B}} h(B^{'})$, it is easy to see that $\pi_{\beta_T}(B)$ increases with T for large T (can be seen by considering derivative of $\pi_\beta(B)$ w.r.t. β). For all other configurations B, $\pi_{\beta_T}(B)$ decreases with T for large T. Hence, $\sum_{T=0}^{\infty}\sum_{B\in\mathcal{B}} |\pi_{\beta_{T+1}}(B) - \pi_{\beta_T}(B)| < \infty$.

Hence, by [16, Chapter 6, Theorem 8.3], $\{V(t)\}_{t\geq 0}$ is strongly ergodic. The expression for the limiting distribution is straightforward to derive. ■

Theorem 4: Under Assumption 1, Algorithm 3 for virtual cache update and Algorithm 2 for real cache update, we have:

$$\lim_{T\to\infty} \frac{\int_0^T \mathbf{P}(R(\tau) = \arg\max_{B\in\mathcal{B}} h(B))d\tau}{T} = 1$$

and hence,

$$\lim_{T\to\infty} \frac{\int_0^T \mathbf{E}(h(R(\tau)))d\tau}{T} = \max_{B\in\mathcal{B}} h(B)$$

Proof: The first part of the proof follows using similar arguments as in the proof of Theorem 2. The second part follows from the first part using the fact that $\mathbf{E}(h(R(\tau))) = \sum_{B\in\mathcal{B}} \mathbf{P}(R(\tau) = B)h(B)$. ■

Remark 5: From [2, Figure 3], we notice that independent placement of contents across BSs can significantly outperform the placement of K most popular contents in each BS cache (for a Poisson distributed network). However, our proposed scheme yields the optimal hit rate for every realization of the location of BSs, so long as the number of BSs is finite. Hence, we can safely claim that our proposed scheme significantly outperforms the placement of K most popular contents in each BS cache.

A. Convergence rate of the virtual cache update scheme

While we are not aware of any closed-form bound on the convergence rate for Algorithm 3, by using [16, Chapter 6, Theorem 7.2], we can provide convergence rate guarantee for Algorithm 1. Let us consider the Markov chain $Y(l) := \{V(lN), V(lN+1), \cdots, V(lN+N-1)\}, l \geq 0$ (indexed by the period index l) evolving under Algorithm 1. Then, by [16, Chapter 6, Theorem 7.2], the total variation distance between μ_l (i.e., the probability distribution of $Y(l)$) and the steady state distribution μ (with marginal distribution equal to $\pi_\beta(\cdot)$) is upper bounded as:

$$d_V(\mu_l,\mu) \leq d_V(\mu_0,\mu)(\delta(Q))^l \leq d_V(\mu_0,\mu)\left(1 - \left(\frac{e^{-\beta\Delta}}{N\binom{M}{K}}\right)^N\right)^l$$

where Q is the transition probability matrix (t.p.m.) of the homogeneous Markov chain $Y(l)$ under Algorithm 1. Clearly, the R.H.S. of the above equation increases with β. Hence, under Algorithm 3, we can expect slower convergence rate as time increases. It has to be noted that there is a trade-off between convergence rate and the accuracy of the virtual cache update scheme using Gibbs sampling; higher accuracy (by taking very large β) obviously requires longer time because of slow convergence rate. It also suggests that the rate of convergence decreases with N (provided that other parameters such as β_0 and Δ are fixed).

V. Learning Content Popularities and Cell Topology

In previous sections, we assumed that the content request arrival rates per unit area, $\lambda_1, \lambda_2, \cdots, \lambda_M$, and the areas $|\mathcal{C}(s)|, s \in 2^{\mathcal{N}}$ are known to all BSs. But, in practice, these quantities may not be known apriori, and one has to estimate these quantities over time as new content requests arrive to the system. In this system, we will extend Algorithm 3 to adapt to learning of these quantities.

At time slot t, the BS j_t (uniformly chosen from the set of BSs) chooses its virtual contents in such a way that the probability of choosing network-wide configuration A at time $t, lN \leq t \leq lN + N - 1$ is $\pi_\beta(A|V_{\cdot,-j_t}(t-1))$.

Let us recall the expression for $h_n(A, s)$ from (9). Clearly, if one can estimate $\lambda_i|\mathcal{C}(s)|$ for all possible $(i, s) \in \mathcal{M} \times 2^{\mathcal{N}}$, then one can have an estimate of $h_n(A, s)$. This can be done by estimating the request arrival rate for content i from the region $\mathcal{C}(s)$; this is easy to do because this is a time-homogeneous Poisson process with rate $\lambda_i|\mathcal{C}(s)|$ request per unit time.

Let us assume that each BS k has an estimate $\hat{\theta}(k, i, s, t)$ for $\lambda_i|\mathcal{C}(s)|$ in slot t. This can be done through continuous message exchange among the BSs which observe the content request arrival process over time.

Now we present the virtual cache update algorithm.

Algorithm 4: This algorithm is same as Algorithm 3 except that the estimate $\hat{\theta}(k, i, s, t)$ is used at slot t by BS k, instead of the actual value of $\lambda_i|\mathcal{C}(s)|$. □

Assumption 2: $\lim_{t\to\infty} \hat{\theta}(k, i, s, t) = \lambda_i|\mathcal{C}(s)|$ almost surely for all $k \in \mathcal{N}, i \in \mathcal{M}, s \in 2^{\mathcal{N}}$.

Assumption 3: $\arg\max_{B\in\mathcal{B}} h(B)$ is unique.

Theorem 5: Under Assumption 2, Assumption 3 and Algorithm 4 for virtual cache update, the discrete time non-homogeneous Markov chain $\{V(t)\}_{t\geq 0}$ is strongly ergodic, and the limiting distribution $\pi_{v,\infty}(\cdot)$ satisfies $\pi_{v,\infty}(\arg\max_{B\in\mathcal{B}} h(B)) = 1$.

Proof: Note that, at a given fixed time $t = T$, given the instantaneous value of estimates, the instantaneous transition probability matrix for $\{V(t)\}_{t\geq 0}$ Markov chain, $Q^{(T)}$, will have a stationary probability distribution. Also, if we assume that there exists exactly one configuration in the set $\arg\max_{B\in\mathcal{B}} h(B)$, then we can say that $\lim_{T\to\infty} |Q^{(T)} - Q^*| = 0$, where Q^* has a stationary distribution which assigns probability 1 on $\arg\max_{B\in\mathcal{B}} h(B)$, and Q^* is strongly ergodic (by Theorem 3). Hence, by [16, Chapter 6, Theorem 8.5], the Markov chain of virtual cache configuration $\{V(t)\}_{t\geq 0}$ is strongly ergodic. ■

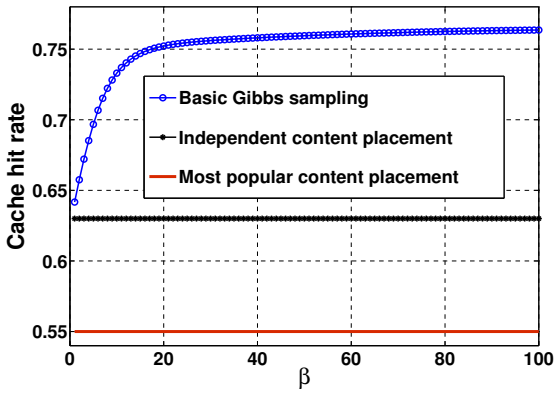

Fig. 2. Comparison of Gibbs sampling based caching strategy, independent content placement strategy, and most popular content placement strategy, for a network with two BSs, two possible contents, and storage capacity for exactly one content in each BS cache. Detailed system parameters can be found in Section VI.

Theorem 6: Under Assumption 1, Assumption 2, Assumption 3, Algorithm 4 for virtual cache update and Algorithm 2 for real cache update, the conclusions of Theorem 4 hold.

Proof: The proof is similar to that of Theorem 4. ■

VI. Performance Improvement Using Gibbs Sampling: A Numerical Example

We consider two BSs placed along a line (say, the x axis). The regions covered by BS 1 and BS 2 are $[0, 6]$ and $[1, 10]$ respectively (note that our theory allows the cells to be asymmetric around the BSs). There are two contents $\mathcal{M} = \{1, 2\}$ with popularities 0.55 and 0.45 respectively. Each BS can store at most one content (i.e., $K = 1$). Content requests are being generated over the region $[0, 10]$ according to a time and space homogeneous Poisson point process with intensity $\frac{1}{10}$ requests per unit time per unit length, and any new request can be for content 1 and content 2 with probabilities 0.55 and 0.45, respectively. The cache hit rate over the region $[0, 10]$ is 0.55 per unit time if the most popular content is placed in both BSs. The cache hit rate is maximized if BS 1 contains content 2 and BS 2 contains content 1, and the cache hit rate is $\frac{1\times 0.45+5\times 1+4\times 0.55}{10} = 0.765$ hits per unit time in this case.

If two BSs choose the contents independently with probabilities r and $(1 - r)$ respectively, the cache hit rate is $0.55r^2 + 0.45(1 - r)^2 + r(1 - r)\frac{1\times 0.55+5\times 1+4\times 0.45}{10} + (1 - r)r\frac{1\times 0.45+5\times 1+4\times 0.55}{10}$, which, upon maximization over r, becomes 0.63 hits per unit time.

If the contents in both caches are chosen probabilistically according to the steady state Gibbs distribution $\pi_\beta(\cdot)$, one can expect that the expected cache hit rate improves as β increases, and converges to the maximum possible cache hit rate (0.765 hits per unit time in our example) as $\beta \uparrow \infty$. We do not consider $\beta = 0$ since it chooses a configuration B uniformly from $\mathcal{B}$.

The above phenomena have been captured in Figure 2. It

also shows that even with finite but large β, significantly higher cache hit rate can be achieved asymptotically compared to the most popular content placement strategy for all BSs, and even w.r.t. independent placement of contents in the BSs. However, if the cells of two BSs do not have any overlap, then placing most popular content will be optimal, and independent placement will be worse than that.

It is to be noted that this example is provided only to demonstrate the potential for performance improvement via Gibbs sampling approach. Providing guarantees for the actual margin of performance improvement for a more realistic network topology (such as Poisson Boolean model for cells) is left for future research endeavours on this topic.

VII. Conclusion

In this paper, we have provided algorithms for cache content update in a cellular network, motivated by Gibbs sampling techniques. The algorithms were shown to converge asymptotically to the optimal content placement in the caches. It turns out that the computation and communication cost is affordable for practical cellular network base stations.

While the current paper solves an important problem, there are still possibilities for numerous interesting extensions: (i) We assumed uniform download cost from the backhaul network for all base stations. However, this is not in general true. Depending on the backhaul architecture, backhaul link capacities and congestion scenario, it might be more desirable to avoid download from some specific base stations. Even different base stations might have different link capacities, and in practice, this will result in queueing delay for the download process. Contents might be of various classes, and hence may not have fixed size. Hence, a combined formulation of cache update and backhaul network state evolution will be necessary. (ii) Different cells might witness different content popularities, but this has not been addressed in the current paper. (iii) Once a content becomes irrelevant (e.g., a news video), it has to be removed completely from all caches; one needs to develop techniques to detect when to remove a content from all caches. (iv) Providing convergence rate guarantees when the inverse temperature is increasing and when arrival rates and cell topology are learnt over time, is a very challenging problem. We leave these issues for future research endeavours on this topic.

Appendix A
Definition of weak and strong ergodicity

Let us consider a discrete-time inhomogeneous Markov chain $\{X(t)\}_{t\geq 0}$ whose transition probability matrix (t.p.m.) between $t = m$ and $t = m + n$ is given by $P(m; n)$. Let $\mathcal{D}$ be the collection of all possible distributions (each element in $\mathcal{D}$ is assumed to be a row vector) on the state space. Then $\{X(t)\}_{t\geq 0}$ is called weakly ergodic if, for all $m \geq 0$,

$$\lim_{n\uparrow\infty} \sup_{\mu,\nu\in\mathcal{D}} d_V(\mu P(m; n), \nu P(m; n)) = 0$$

where $d_V(\cdot, \cdot)$ is the total variation distance between two distributions.

$\{X(t)\}_{t\geq 0}$ is called strongly ergodic if there exists $\pi \in \mathcal{D}$ such that, for all $m \geq 0$,

$$\lim_{n\uparrow\infty} \sup_{\mu\in\mathcal{D}} d_V(\mu^T P(m; n), \pi) = 0$$

References

[1] B. Blaszczyszyn and H.P. Keeler. Studying the sinr process of the typical user in poisson networks using its factorial moment measures. *IEEE Transactions on Information Theory*, 61(12):6774–6794, 2015.

[2] B. Blaszczyszyn and A. Giovanidis. Optimal geographic caching in cellular networks. In *IEEE International Conference on Communications (ICC)*, pages 3358–3363, 2015.

[3] K. Shanmugam, N. Golrezaei, A. G. Dimakis, A. F. Molisch, and G. Caire. Femtocaching: Wireless content delivery through distributed caching helpers. *IEEE Transactions on Information Theory*, 59(12):8402–8413, 2013.

[4] X. Wang, M. Chen, T. Taleb, A. Ksentini, and V. Leung. Cache in the air: Exploiting content caching and delivery techniques for 5g systems. *IEEE Communications Magazine*, 52(2):131–139, 2014.

[5] V. Martina, M. Garetto, and E. Leonardi. A unified approach to the performance analysis of caching systems. In *IEEE Conference on Computer Communications (INFOCOM)*, pages 2040–2048, 2014.

[6] K. Poularakis and L. Tassiulas. Exploiting user mobility for wireless content delivery. In *International Symposium on Information Theory (ISIT)*, pages 1017–1021, 2013.

[7] B. Serbetci and J. Goseling. On optimal geographical caching in heterogeneous cellular networks. *http://arxiv.org/abs/1601.07322*, 2014.

[8] D. Liu and C. Yang. Caching policy toward maximal success probability and area spectral efficiency of cache-enabled hetnets. *https://arxiv.org/abs/1608.03749*, 2014.

[9] K. Avrachenkov, X. Bai, and J. Goseling. Optimization of caching devices with geometric constraints. *https://arxiv.org/abs/1602.03635*, 2016.

[10] K. Hamidouche, W. Saad, M. Debbah, and H.V. Poor. Mean-field games for distributed caching in ultra-dense small cell networks. In *American Control Conference (ACC)*, pages 4699–4704, 2016.

[11] Naveen K.P., L. Massoulie, E. Baccelli, A. Carneiro Viana, and D. Towsley. On the interaction between content caching and request assignment in cellular cache networks. *https://hal.inria.fr/hal-01138204*, 2015.

[12] B.N. Bharath, K.G. Nagananda, and H. V. Poor. A learning-based approach to caching in heterogenous small cell networks. *http://arxiv.org/abs/1508.03517*, 2016.

[13] M. Leconte, G. Paschos, L. Gkatzikis, M. Draief, S. Vassilaras, and S. Chouvardas. Placing dynamic content in caches with small population. *http://arxiv.org/abs/1601.03926*, 2016.

[14] S. Moharir, J. Ghaderi, S. Sanghavi, and S. Shakkottai. Serving content with unknown demand:the high-dimensional regime. *https://arxiv.org/abs/1412.6463*, 2016.

[15] R.M. Karp. Reducibility among combinatorial problems. *Complexity of computer computations*, pages 85–103, 1972.

[16] P. Bremaud. *Markov Chains, Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer, 1999.

Arpan Chattopadhyay obtained his B.E. in Electronics and Telecommunication Engineering from Jadavpur University, Kolkata, India in the year 2008, and M.E. and Ph.D in Telecommunication Engineering from Indian Institute of Science, Bangalore, India in the year 2010 and 2015, respectively. He is currently working in INRIA, Paris as a postdoctoral researcher. His research interests include networks, machine learning, information theory and control.

Bartlomiej Blaszczyszyn received his PhD degree and Habilitation qualification in applied mathematics from University of Wrocaw (Poland) in 1995 and 2008, respectively. He is now a Senior Researcher at Inria (France), and a member of the Computer Science Department of Ecole Normale Suprieure in Paris. His professional interests are in applied probability, in particular in stochastic modeling and performance evaluation of communication networks. He coauthored several publications on this subject in major international journals and conferences, as well as a two-volume book on *Stochastic Geometry and Wireless Networks* NoW Publishers, jointly with F. Baccelli.